

AI 812: Probabilistic Graphical Models

Textbook

- Probabilistic Graphical Models: Principles and Techniques by Daphne Koller and Nir Friedman. MIT Press.
- Handouts and research articles may also be used by the instructor.

Objective

Probabilistic graphical models are a powerful framework for representing complex domains using probability distributions, with numerous applications in machine learning, computer vision, natural language processing and computational biology. Graphical models bring together graph theory and probability theory, and provide a flexible framework for modeling large collections of random variables with complex interactions. This course will provide a comprehensive survey of the topic, introducing the key formalisms and main techniques used to construct them, make predictions, and support decision-making under uncertainty.

Pre-Requisite

Basic probability theory, statistics, programming, algorithm design and analysis.

Course Outcome

During the course, students will learn to develop the knowledge and skills necessary to design, implement and apply these models to solve real problems. Be able to describe directed and undirected graphical models, exact and approximate inference methods and estimation of parameters.

Course Outline

Topics	Allocated Periods
<ul style="list-style-type: none">• Introduction, Probability Theory, Bayesian Networks• Undirected models• Learning Bayes Nets• Exact Inference; Message Passing• Sampling• MAP Inference; Structured prediction• Parameter Learning• Bayesian Learning; Structure Learning• Exponential families; variational inference	45